

Troubleshooting Oracle Performance

Christian Antognini



Apress

Contents

Foreword by Jonathan Lewis	xix
Foreword by Cary Millsap.....	xxi
Foreword from the First Edition.....	xxiii
About the Author	xxv
About the Technical Reviewers	xxvii
Acknowledgments	xxix
Acknowledgments from the First Edition	xxxi
Introduction	xxxiii
■Part I: Foundations	1
■Chapter 1: Performance Problems	3
Do You Need to Plan Performance?.....	3
Requirements Analysis	4
Analysis and Design	5
Coding and Unit Testing	6
Integration and Acceptance Testing.....	7
Designing for Performance	8
Lack of Logical Database Design.....	8
Implementing Generic Tables	8
Not Using Constraints to Enforce Data Integrity	9
Lack of Physical Database Design.....	9
Not Choosing the Right Data Type	9
Not Using Bind Variables Correctly	10
Not Using Advanced Database Features.....	10

CONTENTS

Not Using PL/SQL for Data-Centric Processing	11
Performing Unnecessary Commits	11
Steadily Opening and Closing Database Connections	11
Do You Have Performance Problems?	11
System Monitoring	12
Response-Time Monitoring	12
Compulsive Tuning Disorder	12
How Do You Approach Performance Problems?	13
Business Perspective vs. System Perspective	13
Cataloging the Problems	14
Working the Problems	14
On to Chapter 2	17
■ Chapter 2: Key Concepts	19
Selectivity and Cardinality	19
What Is a Cursor?	21
Life Cycle of a Cursor	21
How Parsing Works	24
Shareable Cursors	25
Bind Variables	29
Reading and Writing Blocks	40
Instrumentation	42
Application Code	43
Database Calls	44
On to Part 2	48
■ Part II: Identification	49
■ Chapter 3: Analysis of Reproducible Problems	51
Tracing Database Calls	51
SQL Trace	52
Structure of the Trace Files	65
Using TRCSESS	67

<i>Profilers</i>	67
<i>Using TKPROF</i>	68
<i>Using TVD\$XTAT</i>	79
Profiling PL/SQL Code	89
<i>Using DBMS_HPROF</i>	89
<i>Using DBMS_PROFILER</i>	96
<i>Triggering the Profilers</i>	101
On to Chapter 4	101
Chapter 4: Real-Time Analysis of Irreproducible Problems	103
<i>Analysis Roadmap</i>	104
<i>Dynamic Performance Views</i>	105
<i>OS Statistics</i>	106
<i>Time Model Statistics</i>	107
<i>Wait Classes and Wait Events</i>	109
<i>System and Session Statistics</i>	113
<i>Metrics</i>	114
<i>Current Sessions Status</i>	115
<i>Active Session History</i>	117
<i>SQL Statement Statistics</i>	126
<i>Real-time Monitoring</i>	126
<i>Analysis With Diagnostics and Tuning Pack</i>	131
<i>Database Server Load</i>	131
<i>System Level Analysis</i>	132
<i>Session Level Analysis</i>	137
<i>SQL Statement Information</i>	139
<i>Analysis Without Diagnostics Pack</i>	142
<i>Database Server Load</i>	142
<i>System Level Analysis</i>	143
<i>Session Level Analysis</i>	147
<i>SQL Statement Information</i>	147
On to Chapter 5	149

Chapter 5: Postmortem Analysis of Irreproducible Problems	151
Repositories	151
Automatic Workload Repository	152
Performing Configuration	152
Taking Snapshots	154
Managing Baselines	154
Statspack	156
Performing Installation	157
Configuring the Repository	157
Taking and Purging Snapshots	158
Managing Baselines	160
Analysis With Diagnostics Pack.....	160
Analysis Without Diagnostics Pack	161
On to Part 3	166
Part III: Query Optimizer.....	167
Chapter 6: Introducing the Query Optimizer.....	169
Fundamentals.....	169
Architecture.....	172
Query Transformations	173
Count Transformation	174
Common Sub-Expression Elimination	174
Or Expansion.....	175
View Merging.....	176
Select List Pruning.....	177
Predicate Push Down	178
Predicate Move Around.....	180
Distinct Placement.....	181
Distinct Elimination.....	181
Group-by Placement.....	181
Order-By Elimination.....	182

Subquery Unnesting	183
Subquery Coalescing.....	183
Subquery Removal Using Window Functions	184
Join Elimination	185
Join Factorization	185
Outer Join to Inner Join	186
Full Outer Join	187
Table Expansion.....	187
Set to Join Conversion.....	188
Star Transformation.....	189
Query Rewrite with Materialized Views.....	189
On to Chapter 7	190
■ Chapter 7: System Statistics	191
The dbms_stats Package	191
What System Statistics Are Available?.....	192
Gathering System Statistics	194
Noworkload Statistics.....	194
Workload Statistics.....	195
Choosing Between Noworkload Statistics and Workload Statistics	198
Restoring System Statistics	199
Working with a Backup Table	200
Logging of Management Operations	200
Impact on the Query Optimizer	202
On to Chapter 8	206
■ Chapter 8: Object Statistics	207
The dbms_stats Package	207
What Object Statistics Are Available?.....	209
Table Statistics	210
Column Statistics.....	211
Histograms	213

Extended Statistics	226
Index Statistics	230
Statistics for Partitioned Objects	232
Gathering Object Statistics	233
Target Objects.....	235
Gathering Options.....	239
Backup Table	244
Configuring the dbms_stats Package	245
The Legacy Way.....	245
The Contemporary Way.....	246
Working with Global Temporary Tables	248
Working with Pending Object Statistics	250
Working with Partitioned Objects.....	251
Challenges	251
Incremental Statistics.....	254
Copying Statistics.....	256
Scheduling Object Statistics Gathering	257
The 10g Way	257
The 11g and 12c Way	259
Restoring Object Statistics	261
Locking Object Statistics.....	262
Comparing Object Statistics	265
Deleting Object Statistics	267
Exporting, Importing, Getting, and Setting Object Statistics.....	268
Logging of Management Operations	269
Strategies for Keeping Object Statistics Up-to-Date	270
On to Chapter 9	271

■ Chapter 9: Configuring the Query Optimizer.....	273
To Configure or Not to Configure.....	273
Configuration Road Map.....	274
Set the Right Parameter!.....	276
Query Optimizer Parameters	276
PGA Management.....	291
On to Chapter 10	297
■ Chapter 10: Execution Plans.....	299
Obtaining Execution Plans.....	299
The EXPLAIN PLAN Statement.....	299
Dynamic Performance Views.....	303
Automatic Workload Repository and Statspack.....	305
Tracing Facilities.....	307
The dbms_xplan Package	311
Output.....	311
The display Function.....	316
The display_cursor Function	322
The display_awr Function	323
Interpreting Execution Plans	325
Parent-Child Relationship.....	325
Types of Operations	328
Stand-Alone Operations.....	328
Iterative Operations	331
Unrelated-Combine Operations	331
Related-Combine Operations.....	333
Divide and Conquer	342
Special Cases	345
Adaptive Execution Plans	348

CONTENTS

Recognizing Inefficient Execution Plans	353
Wrong Estimations.....	353
Restriction Not Recognized	355
On to Part 4	356
Part IV: Optimization.....	357
Chapter 11: SQL Optimization Techniques.....	359
Altering the Access Structures	360
How It Works.....	360
When to Use It	361
Pitfalls and Fallacies.....	361
Altering the SQL Statement.....	361
How It Works.....	361
When to Use It	363
Pitfalls and Fallacies.....	363
Hints	363
How It Works.....	363
When to Use It	370
Pitfalls and Fallacies.....	370
Altering the Execution Environment.....	372
How It Works.....	372
When to Use It	375
Pitfalls and Fallacies.....	375
Stored Outlines.....	375
How It Works.....	375
When to Use It	385
Pitfalls and Fallacies.....	385
SQL Profiles	387
How It Works.....	387
When to Use It	401
Pitfalls and Fallacies.....	402

SQL Plan Management	402
How It Works.....	403
When to Use It	417
Pitfalls and Fallacies.....	417
On to Chapter 12	418
■ Chapter 12: Parsing	419
Identifying Parsing Problems	419
Quick Parses.....	420
Long Parses.....	425
Solving Parsing Problems	427
Quick Parses.....	427
Long Parses.....	433
Working Around Parsing Problems.....	433
Cursor Sharing.....	434
Server-Side Statement Caching	436
Using Application Programming Interfaces	438
PL/SQL.....	439
OCI	442
JDBC.....	443
ODP.NET	445
PHP	446
On to Chapter 13	447
■ Chapter 13: Optimizing Data Access	449
Identifying Suboptimal Access Paths	449
Identification.....	449
Pitfalls.....	452
Causes	454
Solutions.....	454

SQL Statements with Weak Selectivity.....	459
Full Table Scans.....	459
Full Partition Scans.....	461
Range Partitioning	461
Hash and List Partitioning.....	474
Composite Partitioning	475
Design Considerations.....	477
Full Index Scans	479
SQL Statements with Strong Selectivity.....	482
Rowid Access	482
Index Access.....	484
Single-table Hash Cluster Access.....	526
On to Chapter 14	528
Chapter 14: Optimizing Joins	529
Definitions	529
Join Trees	529
Types of Joins.....	534
Restrictions vs. Join Conditions	537
Nested Loops Joins	538
Concept.....	538
Two-Table Join.....	539
Four-Table Join	541
Buffer Cache Prefetches.....	542
Merge Joins.....	544
Concept.....	544
Two-Table Join.....	545
Four-Table Join	548
Work Areas	549
Hash Joins.....	555
Concept.....	555
Two-table Joins	556

Four-Table Joins	557
Work Areas	559
Index Joins	560
Outer Joins	561
Choosing the Join Method.....	562
First-Rows Optimization	562
All-Rows Optimization	562
Supported Join Methods.....	562
Parallel Joins	563
Partition-wise Joins	563
Full Partition-wise Joins.....	563
Partial Partition-wise Joins	566
Star Transformation.....	568
On to Chapter 15	574
Chapter 15: Beyond Data Access and Join Optimization.....	575
Materialized View	575
How It Works.....	576
When to Use It	596
Pitfalls and Fallacies.....	596
Result Caching	597
How It Works.....	597
When to Use It	604
Pitfalls and Fallacies.....	604
Parallel Processing.....	605
How It Works.....	605
When to Use It	636
Pitfalls and Fallacies.....	636
Direct-Path Insert.....	641
How It Works.....	641
When to Use It	644
Pitfalls and Fallacies.....	644

CONTENTS

Row Prefetching	645
How It Works.....	645
When to Use It	650
Pitfalls and Fallacies.....	650
Array Interface.....	650
How It Works.....	650
When to Use It	654
Pitfalls and Fallacies.....	654
On to Chapter 16	654
■ Chapter 16: Optimizing the Physical Design	655
Optimal Column Order	655
Optimal Datatype.....	657
Pitfalls in Datatype Selection.....	657
Best Practices in Datatype Selection.....	660
Row Migration and Row Chaining	662
Migration vs. Chaining.....	662
Problem Description	664
Problem Identification	664
Solutions.....	665
Block Contention	665
Problem Description	666
Problem Identification	666
Solutions.....	670
Data Compression	673
Concept.....	673
Requirements	675
Methods.....	675
■ Part V: Appendix	679
■ Bibliography	681
Index.....	687